

### REMARKS

Claims 1-38 are pending and are rejected. Reconsideration and allowance of Claims 1-38 are respectfully requested.

#### Amendments to Specification

Applicant amends paragraph [0093] to correct a clerical error. No new matter is introduced.

#### Claim Rejections under 35 USC §102 over Blake

Claims 1, 9-10, 22-26, 31-32, 35, and 38 are rejected under 35 USC §102(b) as being anticipated by Blake et al, "An Architecture for Differentiated Services," RFC 2475, December 1998, hereinafter referred to as Blake. Applicant respectfully traverses these rejections.

#### Independent Claim 1

Applicant's Claim 1 recites:

A traffic management processor for processing a plurality of different traffic flows on a per-flow basis, each traffic flow including any number of packets each including a flow identification (ID) indicating to which traffic flow the packet belongs, comprising:

- means for tracking each packet according to its flow ID; and
- means for scheduling each packet according to its flow ID.

As discussed in Applicant's previous response, Blake discloses a traffic routing architecture that processes **aggregated traffic** according to traffic **type** by applying different per-hop behaviors to different traffic class aggregates, where a packets traffic type is indicated by the DS field in the packet headers.<sup>1</sup> More specifically, Blake processes the packets by applying a selected set of per-hop behavior rules according to the packet's traffic type, which as well-known in the art is indicated by marking the DS codepoint of each packet header to a particular traffic type or class. The Office

---

<sup>1</sup> See Blake, page 3, second paragraph.

Action has failed to point to any language in Blake that discloses “means for tracking each packet according to its flow ID” and “means for scheduling each packet according to its flow ID,” as recited in Applicant’s Claim 1.

In rejecting Applicant’s previous arguments, the Office Action states that packets belonging to a particular traffic type are the same as packets belonging to a particular traffic flow, and concludes that the traffic types disclosed in Blake are identical to the traffic flows recited in Applicant’s Claim 1. Applicant disagrees.

As known in the art, a traffic **flow** is NOT the same as a traffic **type** or traffic class. For example, while a traffic flow is made of packets going from a particular source machine to a particular destination machine,<sup>2</sup> a traffic type is an aggregate that includes all traffic assigned to a designated IPv6 priority or per-hop behavior, **regardless of which flow the packets belong to**.<sup>3</sup>

Applicant notes that the Office Action bases the rejection of Applicant’s Claim 1 on Blake while ignoring language in Blake that contradicts the Office Action’s interpretation of a traffic flow. For example, Blake discloses “applying per-hop behaviors to aggregates of traffic”<sup>4</sup> and teaches that “per-application flow or per-customer forwarding state need not be maintained within the core of the network.”<sup>5</sup> Therefore, in contrast to the Office Action’s conclusion that traffic type is “identical” to traffic class, Blake not only distinguishes between traffic flows and traffic types but also seems to teach away from the per-flow tracking and per-flow scheduling recited in Applicant’s Claim 1. Indeed, Blake specifically states that its architecture “should avoid per-microflow or per-customer state” and “should utilize only aggregated classification state.”<sup>6</sup> Further, as known in the art and disclosed in Blake, a packet’s DS codepoint is used to select a per-hop behavior (PHB) for the packet,<sup>7</sup> and thus indicates which class of service the packet is subject to. This is in contrast with the flow ID recited in Claim 1, which indicates which traffic **flow** the packet belongs to.

---

<sup>2</sup> See “IP Quality of Service,” by Srinivas Vegesgan, CCIE, Cisco Press, 2001: An individual traffic flow is made of packets going from an application on a source machine to an application on a destination machine.

<sup>3</sup> See Applicant’s Specification, paragraph [0007].

<sup>4</sup> See Blake, page 3, second paragraph.

<sup>5</sup> See Blake, page 3, second paragraph.

<sup>6</sup> See Blake, page 9, 5<sup>th</sup> and 6<sup>th</sup> bullets.

The Office Action notes that claims should be given the broadest reasonable interpretation. More specifically, the MPEP states that "claims are not to be read in a vacuum, and limitations therein are to be interpreted in light of the specification in giving them their broadest reasonable interpretation."<sup>8</sup> Thus, the words of the claims must be given their plain meaning, as determined by the specification and the prior art. Applicant's specification specifically states that "similar types of traffic are typically aggregated to implement traffic shaping functions. However, aggregating similar types of traffic for queuing and shaping functions does not allow the application of certain QOS and shaping parameters to individual traffic flows."<sup>9</sup> Thus, because both Applicant's specification and Blake distinguish between a traffic flow and a traffic class or type, the traffic flow recited in Applicant's Claim 1 cannot be interpreted to be the same as a traffic type. Accordingly, the flow ID recited in Applicant's Claim 1 cannot be properly interpreted to be the same as the DC codepoint in Blake because a traffic flow is different from a traffic class.

To anticipate a claim under 35 USC §102, each and every element of the claim must be disclosed in a single reference<sup>10</sup>. The exclusion of a claimed element from a prior art reference is typically enough to negate anticipation under 35 USC §102. Thus, because Blake fails to disclose or suggest a traffic management processor including "means for tracking each packet according to its flow ID" and "means for scheduling each packet according to its flow ID," as recited in Applicant's Claim 1, Claim 1 is not anticipated by Blake. Accordingly, Applicant respectfully requests the Office to withdraw the rejection of Claim 1.

Claims 2-11 and 35 depend from Claim 1 and therefore distinguish over the cited references for at least the same reasons as Claim 1.

---

7 See Blake, page 4 (definition of "DS codepoint").

8 See MPEP 2111.01

9 Applicant's specification, paragraph [0007].

10 Corning Glass Works v. Sumitomo Electric, 9 USPQ2d 1962, 1965 (Fed. Cir. 1989).

### Independent Claim 22

Applicant's Claim 22 recites:

A method for processing a number of traffic flows on a per-flow basis, each including one or more packets, comprising:  
receiving an incoming packet;  
determining which traffic flow the incoming packet belongs to; and  
scheduling the incoming packet for departure according to which traffic flow the packet belongs.

Blake does not disclose or suggest a method for processing a number of traffic flows on a per-flow basis, as recited in Applicant's Claim 22.

As discussed above with respect to Claim 1, Blake does NOT disclose or suggest "means for tracking each packet according to its flow ID" or "means for scheduling each packet according to its flow ID," but rather discloses processing **aggregated traffic** by applying per-hop behaviors to packets according to their type using the DS field in the packet headers. Thus, Blake fails to disclose or suggest "determining which traffic flow the incoming packet belongs to," as recited in Applicant's Claim 22, nor has the Office Action pointed to any such language in Blake. Accordingly, Claim 22 is patentable over Blake.

Claims 23-34 and 38 depend from Claim 22 and therefore distinguish over the cited references for at least the same reasons as Claim 22.

### Claim Rejections under 35 USC §102 over Ohgane

Claims 12-21 and 36-37 are rejected under 35 USC §102(b) as being anticipated by Ohgane (USP 5,875,173). Applicant respectfully traverses these rejections.

### Independent Claim 12

Applicant's Claim 12 recites:

A traffic management processor for managing a number of traffic flows each including one or more packets, comprising:

a content address memory (CAM) device having a plurality of rows, each row storing a flow identification (ID) for a corresponding packet, the flow ID indicating to which traffic flow the packet belongs;

a departure time table including a plurality of rows, each coupled to a corresponding row of the CAM device and configured to store a departure time for the corresponding packet; and

compare logic having inputs coupled to corresponding rows of the departure time table, the compare logic for comparing the departure times with each other to determine which departure time is the earliest.

First, Ohgane fails to disclose or suggest a CAM device in which each row stores "a flow identification (ID) for a corresponding packet, the flow ID indicating to which traffic flow the packet belongs," as recited in Applicant's Claim 12.

The Office Action seems to equate the virtual channel (VC) in Ohgane with the flow ID recited in Claim 12. However, as known in the art and described in Ohgane, a virtual channel (VC) is a time-multiplexed signal path across an ATM network through which packets are transmitted.<sup>11</sup> Thus, Ohgane's virtual channel is a signal path, NOT a flow ID indicating which traffic flow a packet belongs to. Accordingly, in contrast to the Examiner's assertion, Ohgane fails to disclose or suggest a CAM device in which each row stores "a flow identification (ID) for a corresponding packet, the flow ID indicating to which traffic flow the packet belongs," as recited in Applicant's Claim 12.

Second, Ohgane fails to disclose or suggest "compare logic for comparing the departure times with each other to determine which departure time is the earliest," as recited in Claim 12.

The Office Action seems to equate Ohgane's logic 512, 515, and 516 with the compare logic of Applicant's Claim 12. However, Ohgane teaches that the packet time values are compared with a *separate counter value*, NOT *with each other* as recited in Claim 12. Indeed, Ohgane specifically states that "when [the packet departure] time value (T) matches with a value identified by the counter 50, an address where this time value (T) is stored can be determined as a virtual channel VC for the next cell to be

---

<sup>11</sup> Ohgane, col. 8, lines 63-67.

transmitted.”<sup>12</sup> This is in marked contrast to the compare logic of Claim 12, which compares the departure times **with each other to determine which departure time is the earliest.**”

Applicants note that comparing the departure times with each other to determine which departure time is the earliest (as recited in Claim 12), rather than comparing the departure times to a current time value (as taught by Ohgane), may allow the traffic management processor of Claim 12 to achieve increased throughput. More specifically, Applicants' specification states that comparing the departure times with each other allows for the selection of a packet for departure in response to every compare operation, independent of any current time value, thereby optimizing packet transmission rates by allowing packets to be continually transmitted.<sup>13</sup> In contrast, prior art systems that “compare departure times with a current time value to select packets for departure may experience idle time during which no packets are transmitted if there is not a match between the current time value and any of the packet departure times.”<sup>14</sup> Therefore, Ohgane fails to disclose or suggest a compare logic that compares the departure times with each other to determine which departure time is the earliest.

Accordingly, because Ohgane does not disclose or suggest “a content address memory (CAM) device having a plurality of rows, each row storing a flow identification (ID) for a corresponding packet, the flow ID indicating to which traffic flow the packet belongs” or “compare logic having inputs coupled to corresponding rows of the departure time table, the compare logic for comparing the departure times with each other to determine which departure time is the earliest,” as recited in Applicant's Claim 12, Claim 12 is patentable over Ohgane.

Claims 13-21 and 36-37 depend from Claim 12 and therefore distinguish over the cited references for at least the same reasons as Claim 12.

---

12 Ohgane, col. 8, lines 19-22.

13 Applicants' Specification, paragraph [0062].

14 Applicants' Specification, paragraph [0063].

Claim Rejections under 35 USC §103

Claims 2-8, 27-30, and 33-34 are rejected under 35 USC §103(a) as being unpatentable over Blake in view of Ohgane. Applicant respectfully traverses these rejections.

Claims 2-8 depend from Claim 1 and therefore distinguish over the cited references for at least the same reasons as Claim 1.


Claims 27-30 and 33-34 depend from Claim 22 and therefore distinguish over the cited references for at least the same reasons as Claim 22.

CONCLUSION

In light of the above remarks, it is believed that Claims 1-38 are in condition for allowance and, therefore, a Notice of Allowance of 1-38 is respectfully requested. If the Examiner's next action is other than allowance as requested, the Examiner is requested to call the undersigned at (408) 236-6646.

Respectfully submitted,

8/3/07  
Dated



William L. Paradice III  
Attorney for Applicant  
Reg. No. 38,990